

CLAIMS

What is claimed is:

- 1) A method for selecting a value in a distributed computing system, the method comprising transmitting a proposed value, a vote for the proposed value, a first proposal identifier and a first step identifier, wherein the vote for the proposed value provides a first recipient device sufficient information to determine, based on the vote for the proposed value and its own vote, whether a first quorum of the distributed computing system has selected the proposed value in a first system step identified by the first step identifier.
- 2) The method of claim 1 further comprising: transmitting a proposal for an operational quorum, a vote for the proposal for the operational quorum, a second proposal identifier and a second step identifier, wherein the vote for the proposal for the operational quorum provides a second recipient device sufficient information to determine, based on the vote for the proposal for the operational quorum and its own vote, whether a second quorum of the distributed computing system has selected the proposal for the operational quorum in a second system step identified by the second step identifier; and receiving an indication of the selection of the proposal for the operational quorum.
- 3) The method of claim 2, wherein the second quorum of the distributed computing system comprises a second transmitting device and the second recipient device, wherein

the second transmitting device transmitted the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein the second recipient device is an inexpensive computing device.

4) A method for selecting a value in a distributed computing system, the method comprising receiving a proposed value, a vote for the proposed value, a first proposal identifier and a first step identifier, wherein the vote for the proposed value provides sufficient information to determine whether a first quorum of the distributed computing system has selected the proposed value in a first system step identified by the first step identifier.

5) The method of claim 4 further comprising selecting the proposed value and transmitting, without waiting for additional messages, the selection of the proposed value to a client device that had originally proposed the proposed value, if the first proposal identifier is greater than or equal to a previously responded to proposal identifier.

6) The method of claim 4 further comprising: receiving a proposal for an operational quorum, a vote for the proposal for the operational quorum, a second proposal identifier and a second step identifier, wherein the vote for the proposal for the operational quorum provides sufficient information to determine whether a second quorum of the distributed computing system has selected the proposal for the operational quorum in a second system step identified by the second step identifier; selecting the proposal for the

operational quorum if the second proposal identifier is greater than or equal to a previously responded to proposal identifier; and transmitting an indication of the selection of the proposal for the operational quorum.

7) The method of claim 6, wherein the second quorum of the distributed computing system comprises a second transmitting device and a second recipient device, wherein the second transmitting device transmitted the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein the second recipient device received the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein further the second recipient device is an inexpensive computing device.

8) A computer-readable medium having computer-executable instructions for selecting a value in a distributed computing system, the computer-executable instructions performing steps comprising transmitting a proposed value, a vote for the proposed value, a first proposal identifier and a first step identifier, wherein the vote for the proposed value provides a first recipient device sufficient information to determine, based on the vote for the proposed value and its own vote, whether a first quorum of the distributed computing system has selected the proposed value in a first system step identified by the first step identifier.

- 9) The computer-readable medium of claim 8, wherein the first quorum of the distributed computing system comprises a first transmitting device and the first recipient device, wherein the first transmitting device transmitted the proposed value, the vote for the proposed value, the first proposal identifier and the first step identifier.
- 10) The computer-readable medium of claim 8, wherein the proposed value is a proposed function to be executed by the distributed computing system.
- 11) The computer-readable medium of claim 10 having further computer-executable instructions performing steps comprising receiving, from the first recipient device, a result of an execution of the proposed function by the first recipient device.
- 12) The computer-readable medium of claim 8 having further computer-executable instructions performing steps comprising: transmitting, to a second quorum of devices in the distributed computing system, a suggested next proposal identifier for the first system step; and receiving, from each device in the second quorum of devices in the distributed computing system, a suggested next proposal identifier response, wherein the suggested next proposal identifier response is null if the each device in the second quorum of devices had not previously voted for the first system step, and wherein the suggested next proposal identifier response comprises an indication of a previously voted for value and a previously voted for proposal identifier, corresponding to the first system step, if the each device in the second quorum of devices had previously voted for the first system step.

13) The computer-readable medium of claim 12 having further computer-executable instructions performing steps comprising: selecting, as the first proposal identifier, a greater identifier than any of the previously voted for proposal identifier; and selecting, as the proposed value, one of the previously voted for value.

14) The computer-readable medium of claim 8 having further computer-executable instructions performing steps comprising: transmitting a proposal for an operational quorum, a vote for the proposal for the operational quorum, a second proposal identifier and a second step identifier, wherein the vote for the proposal for the operational quorum provides a second recipient device sufficient information to determine, based on the vote for the proposal for the operational quorum and its own vote, whether a second quorum of the distributed computing system has selected the proposal for the operational quorum in a second system step identified by the second step identifier; and receiving an indication of the selection of the proposal for the operational quorum.

15) The computer-readable medium of claim 14, wherein the second quorum of the distributed computing system comprises a second transmitting device and the second recipient device, wherein the second transmitting device transmitted the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein the second recipient device is an inexpensive computing device.

16) The computer-readable medium of claim 14, wherein the operational quorum comprises the first quorum of the distributed computing system, and wherein the second system step precedes the first system step.

17) A computer-readable medium having computer-executable instructions for selecting a value in a distributed computing system, the computer-executable instructions performing steps comprising receiving a proposed value, a vote for the proposed value, a first proposal identifier and a first step identifier, wherein the vote for the proposed value provides sufficient information to determine whether a first quorum of the distributed computing system has selected the proposed value in a first system step identified by the first step identifier.

18) The computer-readable medium of claim 17, wherein the first quorum of the distributed computing device comprises a first transmitting device and a first recipient device, wherein the first transmitting device transmitted the proposed value, the vote for the proposed value, the first proposal identifier and the first step identifier, and the first recipient device received the proposed value, the vote for the proposed value, the first proposal identifier and the first step identifier.

19) The computer-readable medium of claim 17 having further computer-executable instructions performing steps comprising selecting the proposed value and transmitting the selection of the proposed value if the first proposal identifier is greater than or equal to a previously responded to proposal identifier.

20) The computer-readable medium of claim 17, wherein the proposed value is a proposed function to be executed by the distributed computing system.

21) The computer-readable medium of claim 20 having further computer-executable instructions performing steps comprising executing the proposed function, and transmitting a result of the execution of the proposed function.

22) The computer-readable medium of claim 21 having further computer-executable instructions performing steps comprising transmitting an indication of a selection of the proposed function.

23) The computer-readable medium of claim 17 having further computer-executable instructions performing steps comprising: receiving a suggested next proposal identifier for the first system step; and transmitting a suggested next proposal identifier response, wherein the suggested next proposal identifier response is null if no vote for the first system step was previously made, and wherein the suggested next proposal identifier response comprises an indication of a previously voted for value and a previously voted for proposal identifier, corresponding to the first system step, if a vote for the first system step was previously made.

24) The computer-readable medium of claim 17 having further computer-executable instructions performing steps comprising: receiving a proposal for an operational

quorum, a vote for the proposal for the operational quorum, a second proposal identifier and a second step identifier, wherein the vote for the proposal for the operational quorum provides sufficient information to determine whether a second quorum of the distributed computing system has selected the proposal for the operational quorum in a second system step identified by the second step identifier; selecting the proposal for the operational quorum if the second proposal identifier is greater than or equal to a previously responded to proposal identifier; and transmitting an indication of the selection of the proposal for the operational quorum.

25) The computer-readable medium of claim 24, wherein the second quorum of the distributed computing system comprises a second transmitting device and a second recipient device, wherein the second transmitting device transmitted the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein the second recipient device received the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein further the second recipient device is an inexpensive computing device.

26) The computer-readable medium of claim 24, wherein the operational quorum comprises the first quorum of the distributed computing system, and wherein the second system step precedes the first system step.

27) A computing device in a distributed computing system implementing a distributed fault-tolerant consensus algorithm, the computing device transmitting a proposed value, a vote for the proposed value, a first proposal identifier and a first step identifier to a first recipient computing device, wherein the vote for the proposed value enables the first recipient computing device to determine whether a first quorum of the distributed computing system has selected the proposed value in a first system step identified by the first step identifier.

28) The computing device of claim 27, wherein the first quorum of the distributed computing system comprises the computing device and the first recipient computing device.

29) The computing device of claim 27 further transmitting, to a second quorum of devices in the distributed computing system, a suggested next proposal identifier for the first system step; and receiving, from each device in the second quorum of devices in the distributed computing system, a suggested next proposal identifier response, wherein the suggested next proposal identifier response is null if the each device in the second quorum of devices had not previously voted for the first system step, and wherein the suggested next proposal identifier response comprises an indication of a previously voted for value and a previously voted for proposal identifier, corresponding to the first system step, if the each device in the second quorum of devices had, for the first system step, previously voted.

30) The computing device of claim 29 selecting, as the first proposal identifier, a greater identifier than any of the previously voted for proposal identifier; and selecting, as the proposed value, one of the previously voted for value.

31) The computing device of claim 27 further transmitting a proposal for an operational quorum, a vote for the proposal for the operational quorum, a second proposal identifier and a second step identifier, wherein the vote for the proposal for the operational quorum enables a second recipient computing device to determine whether a second quorum of the distributed computing system has selected the proposal for the operational quorum in a second system step identified by the second step identifier; and receiving an indication of the selection of the proposal for the operational quorum.

32) The computing device of claim 31, wherein the second quorum of the distributed computing system comprises the computing device and the second recipient computing device, and wherein the second recipient computing device is an inexpensive computing device.

33) A computing device in a distributed computing system implementing a distributed fault-tolerant consensus algorithm, the computing device receiving a proposed value, a vote for the proposed value, a first proposal identifier and a first step identifier, wherein the vote for the proposed value enables the computing device to determine whether a first quorum of the distributed computing system has selected the proposed value in a first system step identified by the first step identifier.

34) The computing device of claim 33, wherein the first quorum comprises the computing device and a first transmitting computing device, wherein the first transmitting computing device transmitted the proposed value, the vote for the proposed value, the first proposal identifier and the first step identifier.

35) The computing device of claim 33 further selecting the proposed value and transmitting, without waiting for additional messages, the selection of the proposed value to a client computing device that had originally proposed the proposed value, if the first proposal identifier is greater than or equal to a previously responded to proposal identifier.

36) The computing device of claim 33 further receiving a suggested next proposal identifier for the first system step; and transmitting a suggested next proposal identifier response, wherein the suggested next proposal identifier response is null if no vote for the first system step was previously made, and wherein the suggested next proposal identifier response comprises an indication of a previously voted for value and a previously voted for proposal identifier, corresponding to the first system step, if a vote for the first system step was previously made.

37) The computing device of claim 33 further receiving a proposal for an operational quorum, a vote for the proposal for the operational quorum, a second proposal identifier and a second step identifier, wherein the vote for the proposal for the operational quorum

enables the computing device to determine whether a second quorum of the distributed computing system has selected the proposal for the operational quorum in a second system step identified by the second step identifier; selecting the proposal for the operational quorum if the second proposal identifier is greater than or equal to a previously responded to proposal identifier; and transmitting an indication of the selection of the proposal for the operational quorum.

38) The computing device of claim 37, wherein the second quorum of the distributed computing system comprises the computing device and a second transmitting device, wherein the second transmitting device transmitted the proposal for the operational quorum, the vote for the proposal for the operational quorum, the second proposal identifier, and the second step identifier, and wherein the computing device is an inexpensive computing device.